

On using milliseconds as a measure of network latency

 devblogs.microsoft.com/oldnewthing/20240206-00

February 6, 2024



Raymond Chen

One of the things I do is serve as an API design reviewer, reviewing and providing feedback on all new APIs added to Windows. There was a network property being added that reported the latency of a network connection. One of the other API reviewers put a note on that property asking, “As network technology improves, will millisecond granularity for reporting latency be sufficient, or should we use microseconds or even nanoseconds?”

I was not the team responsible for the new property, but I felt compelled to clarify the situation: “The speed of light is unlikely to improve.”

(I mean, maybe some scientific breakthrough could result in improvements, but that would result in a much more fundamental change to the design of the new property: It would have to be changed from an unsigned to a signed type!)

Bonus chatter: Now, even with current technology, a smaller unit could still be useful, if the other system is so close by that the latency is less than one millisecond. But once you get that close, the variability in latency begins to overwhelm the actual value. And the target audience for this particular network property was multiplayer online games, so anything less than one millisecond is “amazingly great”. You’d be thrilled to have the problem of having to choose among multiple sub-millisecond game partners!